

Primacy Versus Recency in a Quantitative Model: Activity Is the Critical Distinction

Anthony J. Greene,^{1,2} Colin Prepscius,² and William B. Levy²

²University of Virginia, Department of Neurosurgery, Charlottesville, Virginia 22908 USA

Abstract

Behavioral and neurobiological evidence shows that primacy and recency are subserved by memory systems for intermediate- and short-term memory, respectively. A widely accepted explanation of recency is that in short-term memory, new learning overwrites old learning. Primacy is not as well understood, but many hypotheses contend that initial items are better encoded into long-term memory because they have had more opportunity to be rehearsed. A simple, biologically motivated neural network model supports an alternative hypothesis of the distinct processing requirements for primacy and recency given single-trial learning without rehearsal. Simulations of the model exhibit either primacy or recency, but not both simultaneously. The incompatibility of primacy and recency clarifies possible reasons for two neurologically distinct systems. Inhibition, and its control of activity, determines those list items that are acquired and retained. Activity levels that are too low do not provide sufficient connections for learning to occur, while higher activity diminishes capacity. High recurrent inhibition, and progressively diminishing activity, allows acquisition and retention of early items, while later items are never acquired. Conversely, low recurrent inhibition, and the resulting high activity, allows continuous acquisition such that acquisition of later items eventually interferes with the retention of early items.

Introduction

Following a single exposure to learning, recall is better for items at the beginning (primacy) and

end (recency) of a list than for middle items. This familiar U-shaped serial position curve is taken as evidence for two distinct memory systems (Glanzer and Cunitz 1966). By one account (Waugh and Norman 1965), primacy occurs in a system for long-term memory (LTM), which may maintain information indefinitely, and recency occurs in a system for short-term memory (STM), where unrehearsed information is generally lost in as little as 20 seconds. Accordingly, STM maintains the last few learned items, resulting in recency (Craik et al. 1970). Early list items, which have had the most opportunities for rehearsal in STM, have likewise had the greatest chance to be processed into LTM, resulting in primacy (Rundus 1971). However, some have argued (e.g., Crowder 1982) because initial items have no interference from preceding items and final items have no interference from subsequent items, that a single system could achieve both primacy and recency (Melton 1963; Wixted and Ebbeson 1991).

Compelling evidence for two or more distinct memory systems came first from memory impaired patients and later from primate studies. Impairments of the dorsolateral frontal lobes diminish short-term memory performance and disrupt the recency effect (Shallice and Vallar 1990; Milner et al. 1990). Conversely, damage to the medial temporal lobes, which includes the hippocampal formation and surrounding cortices, severely impairs LTM formation (i.e., anterograde amnesia) and also disrupts the primacy effect (Milner 1970). Importantly, disruption of one system has little effect on the performance of the other (Shallice and Vallar 1990; Castro 1995, 1997), indicating that the systems are distinct and largely independent, and that the formation of LTM does not require prior processing in STM.

In addition to anterograde amnesia, damage to the medial temporal lobes results in temporally graded retrograde amnesia indicating that the formation of LTM involves a labile period of recoding and consolidation, which depends upon the medial temporal lobes (for reviews, see Squire 1992;

¹Corresponding author.

Eichenbaum et al. 1994). This stage of memory is referred to as intermediate-term memory (ITM). This third system also contributes to the maintenance of items or events in memory. By combining all these lines of evidence, we are led to conclude that primacy and recency result from the single-trial learning characteristics of distinct ITM and STM systems.

While it is now well established that STM and ITM are subserved by distinct systems, it is still not clear what specific processing differences characterize these systems, nor is it clear whether such differences constitute incompatibilities that might have favored the evolution of separate memory systems (Sherry and Schacter 1987). Here we examine the processing requirements for primacy and recency to understand the differences that engender distinct systems for ITM and STM.

Our general approach is to model the underlying neural processing mechanisms of behaviorally observable phenomena within a minimal, neurobiologically motivated network architecture designed to learn sequences of input patterns. Several considerations follow directly: To maintain biological plausibility without sacrificing comprehensibility, the model uses McCulloch-Pitts type neurons (i.e., fire or not fire) with sparse asymmetric recurrent connections, local associative synaptic modification, and only a few parameters (Figure 1). Additionally, the simulations are done without reliance on algorithmic features such as a rehearsal mechanism and without external control of stimulus weighting, synaptic modification, or activity modulation. Rather, global network behavior is a purely emergent property of local cellular modifications. While all connectionist models propose to understand information processing from patterns of neuronal connections, our model may be more relevant than back-propagation models, which generally use a class of biologically implausible nonlocal learning rules. The model has proven to be computationally and theoretically informative for a wide variety of cognitive and behavioral simulations. For example, the model has helped explain the neural mechanisms underlying several well-documented learning paradigms, such as transverse patterning and transitive inference (Levy 1996; Levy and Wu 1997). Thus, we believe this architecture provides a basis for understanding the most important characteristics of recurrent networks without making the model excessively confusing or computationally costly. For this paper, the model was configured for single trial learning

to examine quantitatively the similarities and differences between networks that exhibit primacy and recency.

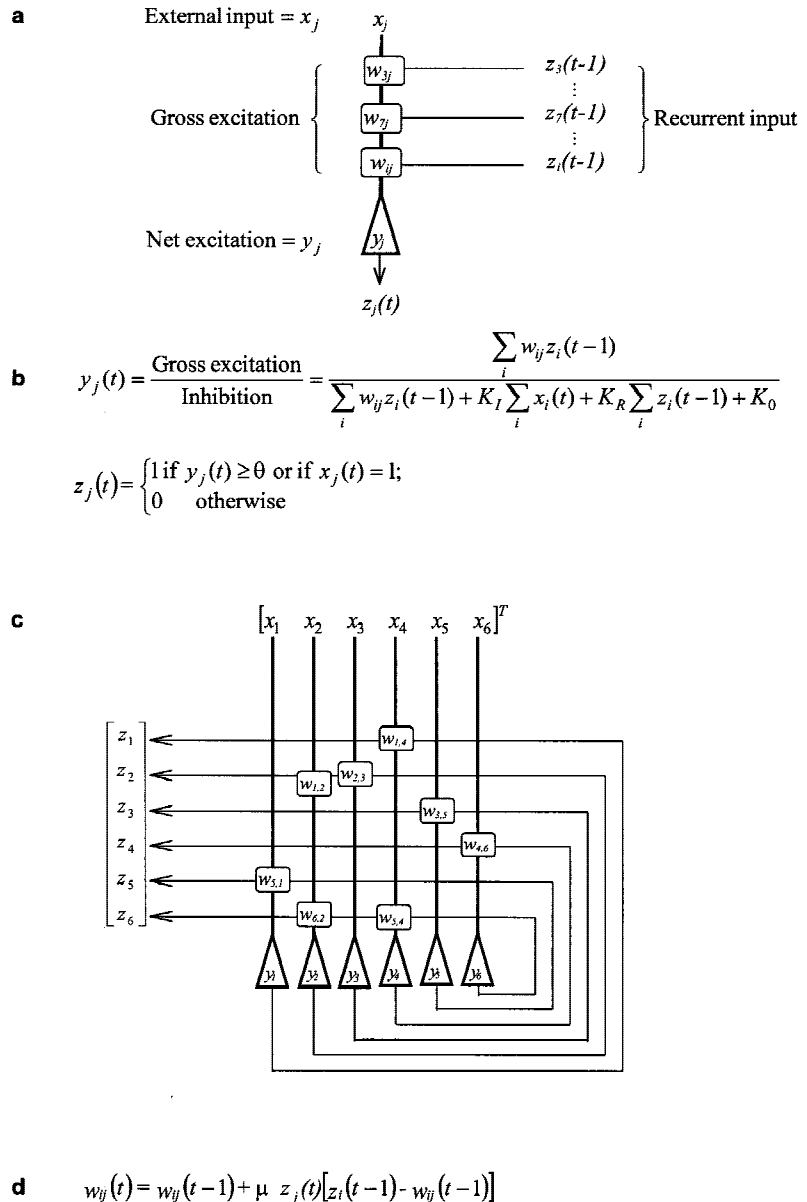
Materials and Methods

Networks were trained on a series of twenty items (each a sequence of neural activations). A high coefficient of synaptic modification (Fig. 1d) allowed single-trial learning (Minai and Levy 1993). Interestingly, networks configured for primacy and recency differed only in the coefficient of recurrent inhibition (K_R controls the inhibition at a given time step based on the total cell firing during the previous timestep; see Fig. 1), which is central to the modulation of activity. Activity is simply the proportion of neurons that fire at any particular time. Activity levels determine the characteristics of learning for reasons we will detail in the Results and Discussion sections. While several factors (Fig. 1) may influence activity levels, recurrent inhibition is a very direct (Minai & Levy 1994) and neurophysiologically well-established (Eccles 1969; Martin 1985; Moser 1996) mechanism of activity control. Furthermore, because recurrent inhibition allows dynamic activity control (i.e., although K_R is fixed, the effect of inhibition varies with activity; see Fig. 1b), activity levels can change throughout training, a property which is axiomatic for serial position effects. A single network of this type could not be configured to simultaneously exhibit both primacy and recency. Consequently these simulations may provide insights that explain why distinct systems are necessary for short and intermediate-term memory.

NETWORK PARAMETERS

Each network simulation incorporated 2000 neurons. Connectivity was 10% (each neuron received connections from exactly 200 other neurons, with no self-connections); for connected neurons, initial synaptic weights, $w_{ij}(0)$, all started at 0.35; firing threshold, 0, was 0.5. In order to achieve single-trial learning, the coefficient of synaptic modification, μ , was 0.8 (for a discussion, see Minai and Levy 1993). The inhibition constants were: $K_I = 0.01$, $K_0 = 0.3$, and $K_R = 0.06$ or $K_R = 0.03$ for primacy- and recency-biased networks, respectively (for a discussion of inhibition and activity control see Smith et al. in press). Appropriate values of K_R were selected empirically by

Figure 1: The neural network model. The model features simple McCulloch-Pitts neurons, shunting inhibition from a single inhibitory interneuron, which loosely controls activity, and sparse, recurrent excitatory connections. (a) A neuron, j , fires at any given timestep if its external input (x_j) is on (i.e., the neuron is forced to fire), or the net recurrent excitation (y_j) meets or exceeds the firing threshold 0. (b) The net excitation of neuron j at time t is the gross excitation of the neuron reduced by inhibitory shunting. The gross excitation is given by the sum of the recurrent inputs ($z_i[t-1]$) multiplied by their respective coefficients of synaptic weights (w_{ij}). K_R is the coefficient of recurrent inhibition, K_I is the coefficient of external inhibition, (consistent with fast-acting inhibitory interneurons (e.g., Buszaki and Eidelberg 1982) K_0 is a constant inhibitory term, corresponding to a resting conductance. (c) Recurrent excitatory connections are sparse and random. (d) Synaptic weights are initially uniform and then are modified according to a postsynaptic modification rule: When the postsynaptic neuron fires, the synaptic weight increases if the presynaptic neuron fired on the previous time step, and it depresses if the presynaptic neuron did not fire. If the postsynaptic neuron does not fire, the synapse is not modified. Although homosynaptic LTD is more widely researched, heterosynaptic modification exists in CA3 and dentate gyrus (Bradler and Barrionuevo 1998; Levy and Steward 1979). Synaptic modification rules for prefrontal cortex are not established. The synaptic modification rate μ controls the speed of learning by scaling the magnitude of synaptic modification.



varying K_R between 0 and 1 in increments of 0.001 and choosing those values which produced the highest overall completion similarity (see below).

ITEMS AND TRAINING

A training epoch consisted of a single presentation of 20 items, each of which constitutes the fundamental learned unit. Each item was a sequence of 10 patterns (or network state vectors), presented in series, one at each timestep t (i.e., vector 1 at timestep 1, vector 2 at timestep 2, etc). Each pattern is analogous to a feature of an item,

such as a phoneme. Note that in behavioral paradigms, stimulus items are known to occupy neural processing time (Burrows and Okada 1975), wherein differing aspects of the stimuli are processed (Cavanagh 1972). An input pattern consisted of 10 externally activated neurons, where external selection guarantees that the neuron will fire (Fig. 1a). At the same time, recurrently activated neurons fire because they have received a sufficient input (above firing threshold) from previously active neurons (Fig. 1a,c). For any training timestep t , the externally and recurrently activated neurons together compose the encoded pattern.

Given that capacity limitations are inherent in any memory system and necessary for serial position effects, we intentionally exceeded the network's capacity so that uniformly good learning was not possible. While all input patterns were orthogonal to one another (i.e., no shared neurons) the encoded patterns could and did overlap. That is, although each neuron was externally activated once and only once, most neurons were also recurrently activated at least once. At each timestep t , the synaptic weights were allowed to vary as governed by the local modification rule (Fig. 1d).

QUANTIFYING OUTPUT QUALITY

Following training, synaptic modification was disabled and an additional set of simulations generated the output used to assess learning. For single-prompt output, we presented the first input pattern (external neurons) from a given item and allowed the network to run without further input to elicit completion of the remaining 9 encoded patterns (external and recurrent neurons) of that item. For full-prompt output, all 10 input patterns of an item were presented in sequence to generate the encoded patterns of that item. Single-prompt output defines the network's recall and full-prompt output defines the correct codes. The single-prompt outputs were then compared to full-prompt outputs for accuracy.

The quality of output is completion similarity, which is a comparison of input and output vectors. Specifically, for a given item, completion similarity was computed by averaging cosine similarities (for details, see Kohonen 1997) between single-prompt output vectors and corresponding full-prompt output vectors. Two qualifications exist: In rare instances, a given pattern's single-prompt output was a closer match to another pattern's full-prompt output than it was to its own, in which case the cosine similarity for that pattern was set to 0 (recall failure). Secondly, the first (prompted) vector for each item was omitted from the computation of completion similarity because the cosine values are always 1.0. To assess the accuracy of a given item, the remaining 9 cosines for each item were averaged together to yield the mean cosine similarity measure for that item. This entire training-testing procedure was repeated 25 times, each with a different random initialization seed (connectivity remained constant at 10% but the choice of connected neurons is randomized, as is the choice of neurons active for the first timestep t). For a given

item the mean cosine similarities from all 25 simulations were averaged to produce the completion similarity for that item. The averages are shown in Figure 2 (standard error for all items is less than 3%).

ASSESSING RETENTION AND ACQUISITION

For retention simulations, training consisted of a single, sequential presentation of all 20 items. Network output was generated and assessed as described above. The completion similarity of each item was graphed against the serial position of that item (Figs. 2a,b). While retention is the usual measure of memory quality, failure to retain information may result when items are never learned or when learned items are forgotten. The distinction between forgetting and failure to acquire is essentially indistinguishable in behavioral and cognitive paradigms, while neural network models are well suited to analyze this distinction.

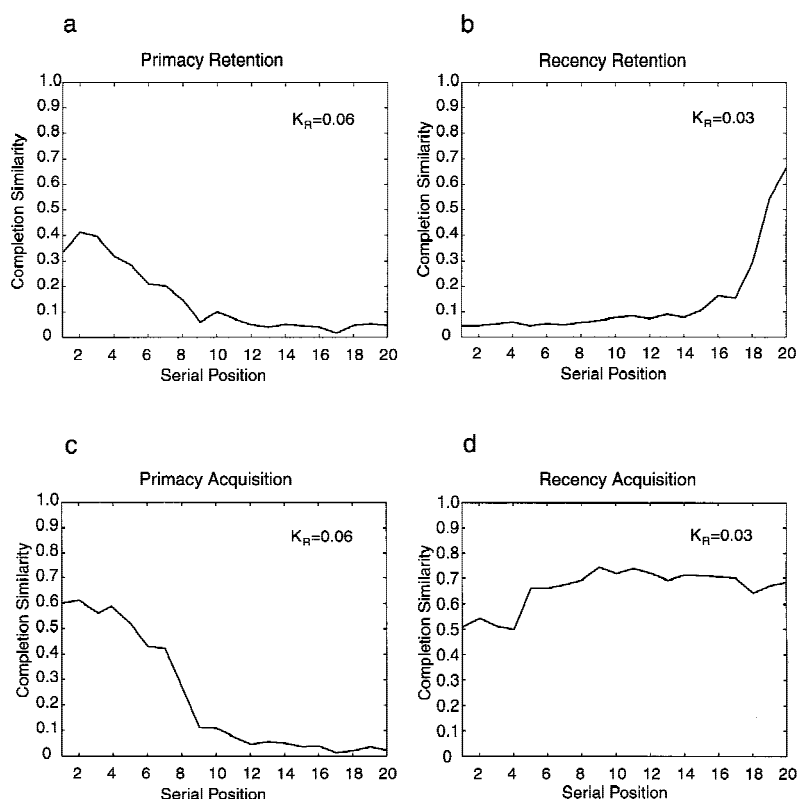
Acquisition was assessed by partial training and then testing. Acquisition simulations for item n consisted of training a network on items 1 through n only, followed by output production and completion similarity computations for only item n . This allowed an assessment of learning without any overwriting from subsequent items. Identically parameterized simulations produced output vectors for all 20 items. For a given parameterization, the set of 20 simulations is equivalent to a single simulation in which acquisition is assessed after each item without disrupting training. Item acquisition for primacy- and recency-biased networks was graphed against serial position (Figs. 2c,d).

ASSESSING THE ACQUISITION-ACTIVITY INTERACTION

An additional set of simulations (Figs. 3a,b) amplify the information available from the acquisition simulations (Figs. 2c,d) by including the activity levels observed during network training and the changes in acquisition associated with varied activity. First, to measure activity during learning, simulations were performed for an entire training epoch for both a primacy- and a recency-biased network. Activity level was recorded for each item. Again, the procedure was repeated for 25 initialization seeds and the output was averaged to produce average item activity for each item (Figs. 3a, b, thick black line).

Because activity level and training experience

Figure 2: Primary versus recency. A network exhibits primacy when parameterized with a relatively high coefficient of recurrent inhibition [$K_R = 0.06$] and recency when parameterized with a lower coefficient of recurrent inhibition [$K_R = 0.03$]. For primacy (a), the earlier items were recalled best during retention tests. For recency (b), later items were recalled better than earlier items during retention tests. In contrast to retention, acquisition testing determines how well an item is learned without interference from subsequent learning. In the primacy case c, early items were acquired whereas later ones were not. Thus, the primacy retention result is explained by the absence of later acquisition where this absence protected earlier learning from overwriting. On the other hand, acquisition assessment indicates that recency simulations (d), are continuously learning and, therefore, earlier items are overwritten as additional items are presented. The completion similarity measure is the average cosine between input and output vectors for all patterns of a given sequence. For retention simulations networks are trained on all 20 sequences, synaptic modification is turned off and the output sequence of vectors for all items are generated. For acquisition simulations, a given item n was assessed by training networks on items 1 through n , turning off synaptic modification and cueing for the output of item n .



change simultaneously throughout training, we devised an analysis to distinguish the effects of these two variables. Functionally, a network learns $n-1$ items exactly as it would in the primacy or recency case, then activity is varied to sample all possible activity levels, and the acquisition performance is measured at each sampled activity level. For fixed values of K_R (0.06 for primacy and 0.03 for recency), and a fixed serial position for an item ($n = 1, 2, \dots$, or 20), a network was trained on items 1 through $n-1$. Then, we varied K_R between 0 and 1 in increments of 0.001 and trained the final item n . The K_R variation cycle was repeated for all n between 1 and 20. Acquisition accuracy and network activity were both assessed across all items and all increments of K_R . Again, this entire procedure was repeated for 25 network configurations, the results of which were averaged together, by item.

The acquisition-activity pairs for each serial position are expressed as contour plots (Figs. 3a,b, colored lines). Data pairs containing 0 activity were omitted. For a given item, data pairs from all simulations were sorted with respect to activity

into a vector of pairs. To generate smooth contour data, each vector was then divided into bins of 25 pairs each, averages were computed for each bin, and finally new pairs of performance (acquisition completion similarity) and activity were interpolated from a quadratic fitted to the averages. Contour plots were graphed from the resulting matrices of points. One can see that the activity levels observed in the networks (Figs. 3a,b, thick black line), determine the probability of acquisition (Figs. 2c,d).

Results and Discussion

Recurrent inhibition, and its control of activity, determined item retention (Figs. 2a,b). Importantly, when recurrent inhibition was varied from lower to higher levels, there was a smooth transition from recency, to uniformly poor learning, to primacy. When parameterized with a high coefficient of recurrent inhibition, the output of the network tended toward primacy with good retention for the first few items, while retention of later

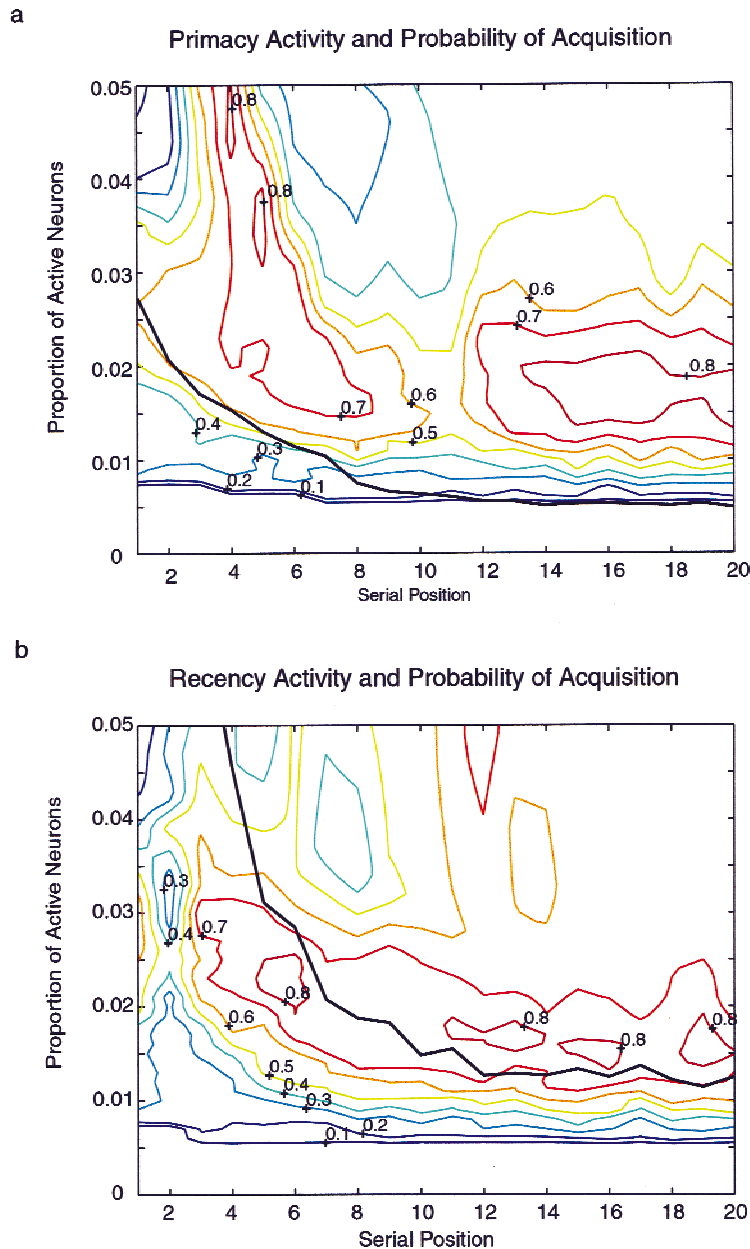


Figure 3: Separating activity effects from serial position effects. These contour plots show the relationship between acquisition and activity for each sequential item. The black line of each figure shows the observed activity levels for each successive item in simulations parameterized for primacy (a) and recency (b). The colored contour lines represent a surface where altitude indicates acquisition quality at various activity levels for each item. Activity was controlled indirectly by iterating through values of K_R . For any item n , networks were trained on items 1 through $n-1$, K_R was altered and item n was trained. Thus, by varying activity on the target trial, we separate the effect of activity level from the effect of serial position. Note how the primacy-biased simulations briefly maintain high enough activity levels to allow early acquisition, but the activity quickly drops to a low activity level where acquisition is highly unlikely. The recency-biased simulations do not drop to so low a level resulting in a moderate to high probability of acquisition throughout training.

items was quite poor. When parameterized with a low coefficient of recurrent inhibition, networks show a clear tendency toward recency with better retention for later items than for early items. The differences between primacy- and recency-biased networks will be shown to depend on how items are differentially acquired, and which of those are retained. The mechanisms of acquisition and retention will be explained as a consequence of activity control.

Acquisition measures how well a given item is learned before either time or the introduction of a

new item can diminish learning. As with the retention simulations, the level of recurrent inhibition determined which items were acquired (Figs. 2c,d). In simulations configured for primacy, acquisition was high for early items but then dropped precipitously, indicating that only early items were learned. In networks configured for recency, acquisition remained fairly constant, which means that continual learning occurred. Taken together, acquisition and retention of items provide a foundation for understanding why primacy and recency occur in these simulations. For primacy, early

items were acquired and retained, and later items were never acquired. For recency, all items were acquired but only the most recent items were retained. The precise mechanisms that allow selective acquisition in primacy-biased networks and selective retention in recency-based networks are a direct consequence of the control of systemic activity.

THE CENTRAL ROLE OF ACTIVITY

The fundamental difference between networks exhibiting primacy and those exhibiting recency is the dynamic modulation of activity, here engendered by recurrent inhibition. The function of inhibition is to govern overall network activity. In particular, the shunting effect of recurrent inhibition increases when more neurons are active because a proportion (K_R) of activity from a given t is fed back as inhibition for the subsequent t (Fig. 1b). To see activity changes, we plotted network activity for both primacy- (higher recurrent inhibition) and recency-biased networks (lower recurrent inhibition) against serial position (Figs. 3a,b, thick black lines). It is apparent that with higher recurrent inhibition (primacy-biased) activity dropped more quickly to low levels, while with lower recurrent inhibition (recency-biased) activity was maintained at a relatively higher overall level.

Network activity determines the characteristics of learning in at least three important ways: First, Hebbian synaptic modification depends directly on cell firing. In predominantly positive feedback networks, the greater the number of active neurons, the greater the likelihood that any given synapse will be modified (Figs. 1c,d). Second, as activity levels decrease, it becomes increasingly unlikely that synapses will exist between sequentially active neurons. Recall that the model employs biologically motivated sparse connectivity such that any neuron was connected to only 10% of other neurons. For example, at the lowest possible activity level for these simulations (only 10 external neurons active), ~35% of the active neurons would be expected to have no available inputs from the previously active neurons ($p = (0.9)^{10} = 0.35$). On the other hand, at higher activity levels external neurons are accompanied by numerous recurrently activated neurons, which are activated solely by their existing connectivity from other active neurons. For these two reasons, as activity decreases, there is a rather sudden drop in the probability of acquisition that results from insufficient

recurrent activity. Third, higher activity levels also mean that the system's capacity will be lower. For example, if we suppose that activity did not change over time, the theoretical capacity for nonoverlapping (orthogonal) patterns would be inversely proportional to the fraction of the neurons active for any one pattern (Levy and Wu 1996). A learned item becomes disrupted if the neurons involved in its representation are later co-opted to become part of the representation of other patterns via associative modification of their input synapses. When acquisition is ongoing, the earliest items are most likely to be lost from memory because they suffer the most overwriting. Simply put, higher activity means that learning is more likely to occur, but it also reduces the overall capacity such that older items tend to be overwritten more quickly by newer items. Interestingly, it is this tendency to overwrite the oldest items that causes recency to occur in high activity networks, and it is also why primacy is observed only in networks in which early learning is protected by diminishing activity.

The relationship between activity and acquisition (Figs. 3a,b, colored contour lines connect points of equal acquisition probability) is central to understanding the differences between primacy- and recency exhibited by these simulations. As can be seen, for both the primacy and recency-biased networks, activity below ~1.0% was too low to support learning. In the primacy-biased network, activity dropped quickly to levels where acquisition was highly unlikely. On the other hand, the recency-biased network maintained activity levels throughout training that corresponded to moderately high acquisition.

To recapitulate, the low recurrent inhibition network (biased towards recency) remained at activity levels that promoted a high probability of acquisition throughout training such that incoming items were continuously learned. But the high activity levels also diminished capacity and increased the likelihood that old items would suffer overwriting. Thus all items were acquired, but only items late in learning were retained. On the other hand, the high recurrent inhibition network (biased toward primacy) began at activity levels that promoted a high probability of acquisition, but activity descended to levels where acquisition was very unlikely. Items encountered early were acquired, well because of high initial activity levels. Declining activity levels did not support synaptic modification or provide enough synapses to encode later

items. Since later items were not acquired, nor were many synapses modified, there was only minimal disruption to the retention of items acquired early in learning.

RELEVANCE TO BEHAVIORAL, COGNITIVE AND BIOLOGICAL LITERATURES

Several considerations are important for comparing this learning model to human and animal learning models. For each simulation, a network is initialized (synaptic weights are reset to initial values) prior to learning. Because a recency-biased network retains the most recent patterns, continued learning without reinitialization should make no difference. Thus, because STM is well-established as an ongoing learning system (Waugh and Norman 1965), the characteristics of the recency-biased network are most relevant after activity settles to asymptotic levels (i.e., after the first ten or so items have been learned). However, this is not the case for a primacy-biased network. If ongoing stimulus presentation is assumed, new learning in a primacy-biased network must be accompanied by a boost in activity (N.B., in our simulations, this activity increase is a by-product of initialization). In fact, fMRI evidence shows that medial temporal lobe activity is temporarily increased when novel stimuli are introduced (Dolan and Fletcher 1997). Additionally, single-unit electrode recordings from the CA1 region of the rat hippocampus reveal that novelty induced learning erases recent synaptic modification via a reversal of previously induced long-term potentiation (Xu et al. 1998). Note that the two critical features of initializing a network in these simulations are an initially high activity level and a randomization of synaptic weights. Such findings are highly consistent with the notion that the introduction of novelty causes the ITM system to reset for new learning. While we do not speculate on where or how novelty is detected, the resultant increase in medial temporal lobe activity and reversal of LTP tend to support our hypothesis of primacy.

Several aspects of this model differ substantially from many current cognitive and behavioral models of STM and ITM. First, a rehearsal mechanism is widely believed to explain the primacy effect. Our model does not achieve primacy by preferentially rehearsing initial items. By most rehearsal accounts, items are intentionally maintained in STM such that initial items are better rehearsed and thereby more likely to be processed

into ITM. There are several problems with rehearsal hypotheses of primacy. Maintenance rehearsal is neither sufficient nor necessary for ITM formation (Craig and Watkins; 1973; Yates and Curley 1986). Additionally, STM impaired patients show that an intact STM system is not necessary for primacy to occur (Shallice and Vallar 1990). Finally, primates show primacy in list learning (Castro and Larsen 1992), but intentional rehearsal strategies are dubious in this case. Models of ITM that do not require rehearsal are not only more parsimonious, they are also more consistent with convergent empirical findings.

A second area of divergence between this model and others is that here dynamic activity modulation (governed by recurrent inhibition) is hypothesized to be the critical difference between systems for STM and ITM. According to most accounts of serial position effects, ITM is a slow-learning system and STM is a fast-learning system (Atkinson and Shiffrin 1968). However, for these simulations, primacy and recency are achieved with identical rates of synaptic modification (Fig. 1d). The hypothesis of slow- and fast-learning systems gained favor when it was believed that the very purpose of STM was to rehearse information for the ITM/LTM system (ITM and LTM had not yet been distinguished). However, because ITM does not require rehearsal in STM and since both STM and ITM are capable of single-trial learning, the hypothesis of slow- and fast-learning systems loses a great deal of support. The architectures we used for primacy- and recency-biased networks were absolutely identical except for changes in the parameterization of recurrent inhibition. The hypothesis put forth here is that recurrent inhibition may constitute a sufficient difference between systems for STM and ITM. Specifically, we do not propose that biological systems for STM and ITM necessarily share identical architectures, but these results do suggest that dynamic modulation of activity may be the critical distinction between such systems.

The present model also provides a neural basis for theories that explain serial position effects in terms of interference (Underwood 1957). By such accounts, learning a given item can prevent the acquisition of future items (proactive interference) or it can disrupt previous learning (retroactive interference). Primacy is believed to be the result of proactive interference (e.g., Han et al. 1998) and recency is the consequence of retroactive interference. Our model extends this theoretical perspective by hypothesizing clear, precise, and internally

consistent mechanisms of interference for differential acquisition and retention. That is, proactive inhibition is exhibited in networks parameterized with high recurrent inhibition such that the activity during early learning drives down subsequent activity to prevent later learning as well as the consequent overwriting of early items. Likewise, retroactive interference is exhibited in networks parameterized with low recurrent inhibition because higher asymptotic activity allows ongoing learning to disrupt early learning.

The effect of activity on acquisition also explains why primacy and recency are not observed simultaneously in a single network. Simply put, activity profiles which protect early learning (primacy) do so at the expense of later learning (recency), and vice-versa. Because all items have the same external activity and are presented only once for learning, without rehearsal, the only differences between learned and unlearned patterns are the levels of activity during learning. A single system of this type could not be made which would allow new items to overwrite old and simultaneously preserve initial items because all acquired items are either vulnerable to overwriting or they are not (c.f., the stability-plasticity problem; Grossberg 1982). Thus in this type of model (with fixed parameters and no rehearsal) primacy and recency constitute mutually exclusive events.

Finally, we offer several straightforward predictions. While engaged in single-trial or novel learning, ITM should exhibit identifiable activity maxima corresponding to better memory for items at the beginning of novel lists, while STM should exhibit a more constant level of activity, corresponding to consistent acquisition of all incoming items. Furthermore, the distinction between STM and ITM should be observable as a stable activity level in STM (recall that the asymptotic activity is argued to be the relevant portion of the recency model) compared to a precipitous drop in activity for ITM. Thus, differences in relative activity across items should correspond to the characteristic acquisition and retention gradients for ITM and STM systems.

Acknowledgments

We thank Ali Minai, Paul Rodriguez, Donald Christman, Sean Polyn, and Aaron Shon for their helpful comments on the manuscript. This work was supported by NIH F32 MH11979 to AJG and by NIH RO1 MH48161 and RO1 MH57358 to WBL.

The publication costs of this article were defrayed in part by payment of page charges. This article must therefore be hereby marked "advertisement" in accordance with 18 USC section 1734 solely to indicate this fact.

References

- Atkinson, R.C. and R.M. Shiffrin. 1968. Human memory: A proposed system and its control processes. In *The psychology of learning and motivation: Advances in research and theory*, vol. 2. (ed. K.W. Spence) pp. 89-195. Academic Press, New York.
- Bradler J.E. and G. Barrionuevo. 1990. Heterosynaptic correlates of long-term potentiation induction in hippocampal CA3 neurons. *Neuroscience* **35**: 265-271.
- Burrows, D. and R. Okada. 1975. Memory retrieval from short and long lists. *Science* **183**: 1031-1033.
- Buzaki, G. and E. Eidenlberg. 1982. Direct afferent excitation and long-term potentiation of hippocampal interneurons. *J. Neurophysiol.* **48**: 597-607.
- Castro, C.A. and T. Larsen. 1992. Primacy and recency effects in nonhuman primates. *J. Exp. Psychol.: Animal Beh. Processes* **18**: 335-340.
- Castro, C.A. 1995. Primacy and recency effects in Rhesus monkeys (*Macaca mulatta*) using a serial probe recognition task: I. Effects of diazepam. *Psychopharmacology* **119**: 421-427.
- . 1997. Primacy and recency effects in Rhesus monkeys (*Macaca mulatta*) using a serial probe recognition task. II Effects of atropine sulfate. *Behav. Neurosci.* **111**: 676-682.
- Cavanagh, J.P. 1972. Relation between the immediate memory span and the memory search rate. *Psychol. Rev.* **79**: 254-257.
- Craik, F.I.M. and M.J. Watkins. 1973. The role of rehearsal in short-term memory. *J. Verbal Learning and Verbal Behavior* **12**: 599-607.
- Craik, F.I.M., J.M. Gardiner, and M.J. Watkins. 1970. Further evidence for a negative recency effect in free recall. *J. Verbal Learning and Verbal Behavior* **9**: 554-560.
- Crowder, R.G. 1982. The demise of short-term memory. *Acta Psychologica* **50**: 291-323.
- Dolan, R.J. and P.C. Fletcher. 1997. Dissociating prefrontal and hippocampal function in episodic memory encoding. *Nature* **388**: 582-585.
- Eccles, J.C. (1969). *The inhibitory pathways of the central nervous system*. Sherrington Lectures, Liverpool University Press, Liverpool.
- Eichenbaum, H., T. Otto, and N.J. Cohen. 1994. Two functional components of the hippocampal memory system. *Behav. Brain Sci.* **17**: 449-518.

- Glanzer, M. and A.R. Cunitz. 1966. Two storage mechanisms in free recall. *J. Verbal Learning and Verbal Behavior*. **5**: 351–360.
- Grossberg, S. 1982. *Studies of mind and brain: Neural principles of learning, perception, development, cognition and motor control*. Boston: Reidel.
- Han, J.S., M. Gallagher, and P. Holland. 1998. Hippocampal lesions enhance configural learning by reducing proactive interference. *Hippocampus* **8**: 138–146.
- Kohonen, T. 1997. In *Self-organizing maps (2nd Ed)*. editor Springer, Berlin.
- Levy, W.B. 1996. A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus* **6**: 579–590.
- Levy, W.B. and O. Steward. 1979. Synapses as associative memory elements in the hippocampal system. *Brain Res.* **17**: 233–245.
- Levy, W.B. and X.B. Wu. 1996. The relationship of local context codes to sequence length memory capacity. In *Network: Computation in Neural Systems* **7**: 371–384.
- . 1997. A simple, biologically motivated network solves the transitive inference problem. *IEEE International Conference on Neural Networks 1997*, **1**: 369–371.
- Martin, J.H. 1985. Cortical neurons, the EEG, and the mechanisms of epilepsy. In *Principles of neural science* 2d Edition. (ed. E.R. Kandel and J.H. Schwartz) pp. 636–647. Elsevier, New York.
- Melton, A.W. 1963. Implications of short-term memory for a general theory of memory. *J. Verbal Learning and Verbal Behavior* **2**: 1–21.
- Milner, B. 1970. Memory and the medial temporal regions of the brain. In *Biology of memory*. (ed. K.H. Pribram and D.E. Broadbent) Academic Press, New York.
- Milner, B., P. Corsi, and G. Leonard. 1991. Frontal-lobe contribution to recency judgments. *Neuropsychology* **29**: 601–618.
- Minaï, A.A. and W.B. Levy. 1993. Sequence learning in a single trial. *INNS world Congress on Neural Networks II*, 505–508.
- . 1994. Setting the activity level in sparse random networks. *Neural Computation* **6**: 83–97.
- Moser, E.I. 1996. Altered inhibition of dentate granule cells during spatial learning in an exploration task. *J. Neurosci.* **16**: 1247–1259.
- Rundus, D. 1971. Analysis of rehearsal processes in free recall. *J. Exp. Psychol.* **89**: 63–77.
- Shallice, T. and G. Vallar. 1990. The impairment of auditory-verbal short-term storage. In *Neuropsychological impairments of short-term memory*. (ed. G. Vallar and T. Shallice) (pp. 187–213). Cambridge University Press, Cambridge.
- Sherry, D.F. and D.L. Schacter. 1987. The evolution of multiple memory systems. *Psychol. Rev.* **94**: 439–454.
- Smith, A.C., X.B. Wu, and W.B. Levy. In Press. Controlling Activity Fluctuations in Large, Sparsely Connected Random Networks, *Network: Computation in Neural Systems*.
- Squire, L.R. 1992. Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.* **99**: 195–231.
- Underwood, B.J. 1957. Interference and forgetting. *Psychol. Rev.* **64**: 49–60.
- Waugh, N.C. and D.A. Norman. 1965. Primary memory. *Psychol. Rev.* **72**: 89–104.
- Wixted, J.T. and E.B. Ebbesen. 1991. On the form of forgetting. *Psychol. Sci.* **2**: 409–415.
- Xu, L., R. Anwyl, and M.J. Rowan. 1998. Spatial exploration induces a persistent reversal of long-term potentiation in rat hippocampus. *Nature* **394**: 891–894.
- Yates, J.F. and S.P. Curley. 1986. Contingency judgment: Primacy effects and attention decrement. *Acta Psychologica* **62**: 293–302.

Received September 7, 1999; accepted in revised form December 2, 1999.